BatAlloc: Effective Battery Allocation against Power Outage for Cellular Base Stations

Fangxin Wang*, Feng Wang†, Xiaoyi Fan*, Jiangchuan Liu*

*Simon Fraser University, †University of Mississippi
fangxinw@sfu.ca,fwang@cs.olemiss.edu,xiaoyif@sfu.ca,jcliu@cs.sfu.ca

ABSTRACT

Base stations play a key role in today's cellular networks. Their reliability and availability heavily depend on the electrical power supply. Modern power grid is known to be highly reliable, but still suffers from outage due to severe weather or human-driven accidents, particularly in remote areas. Most of the base stations are thus equipped with backup battery groups. Given their limited numbers and capacities, they however can hardly sustain a long power outage without a proper allocation strategy. A deep discharge will also accelerate the battery degradation and eventually contribute to a higher battery replacement cost.

In this paper, we closely examine the power outage events and the backup battery status from a one-year dataset of a major cellular service provider, including 4206 base stations distributed across 8400 square kilometers and more than 1.5 million records on battery activities. We then develop *BatAlloc*, a battery allocation framework to address the mismatch between the battery supporting ability and diverse power outage incidents. We build up a deep leaning based approach to accurately profile battery features and present an effective solution that minimizes both service interruption time and the overall cost. Our trace-driven experiments show that BatAlloc cuts down the average service interruption time from 5 hours to nearly zero with only 88% of the overall cost compared to the current practical allocation.

KEYWORDS

Backup power system, Battery feature profiling, Deep learning, Battery allocation

1 INTRODUCTION

Wireless mobile networks, particularly wide-area cellular networks, have seen deep penetration and broad coverage in the past decades. Base stations play a key role in today's cellular networks. Their reliability and availability heavily depend on the electrical power supply, for such modules as transceivers, air conditioners, monitoring system are all power hungry. Modern power grid is known to be highly reliable in urban areas, but still suffers from outage due to severe weather (e.g., storm, heavy rain, hurricane, fire, earthquake) or human-driven accidents (e.g., vandaflism or theft) [3, 9]. In many rural areas, the outage can be quite frequent, no matter in developing or developed countries.

To avoid service interruptions, most base stations are equipped with energy-storage battery groups as the backup power. These batteries are usually kept in the float charge state. When a power outage happens, they will be activated to maintain cellular services until the electrical grid recovers or diesel generators are launched. The capacity of a backup battery group is limited, which typically lasts 10 to 12 hours during power outage. For remote areas or during extreme weather, however, the power recovery can take a long time (e.g., during the severe windstorm in March 2010, the power outage in southwestern Connecticut as well as parts of Long Island and New Jersey lasted way over ten hours, and in some of the rural communities the outage lasted as long as 6 days [1]), so for technicians to arrive at the base station with diesel generators, not to mention that many base stations would be affected at the same time. As such, a long power outage without timely rescue will inevitably drain the backup battery, resulting in service interruption during the extended power outage. This in turn seriously affects the user experience and undermines the telecom operators' service commitments, particularly considering the clients' high reliance on the network during the incident.

Moreover, different from batteries for phones or electrical vehicles which regularly experience full charge/discharge cycles, a deep discharge of an energy-storage battery group (typically lead acid) will severely affect its internal structure, reducing its capacity and lifetime. Given the long time interval between regular maintenances (typically three months [3]), the poor working condition of the battery after a deep discharge will further accelerate its degradation. In the worst case, an overdischarge can permanently damage the battery. Considering the transportation and labor costs, an emergent battery replacement and maintenance can be prohibitively expensive, particularly for remote areas.

In this paper, we closely examine the power outage events and the backup battery activities from a one-year dataset (from July-28-2014 to July-28-2015) of a branch of China Mobile (the biggest cellular service provider in China), including 4206 base stations and more than 1.5 million records on battery activities. These base stations are distributed over an area of 8400 square kilometers with over 10,616,000 clients. Our analysis of the data reveals the ineffectiveness of existing battery allocation strategies during power outages. In particular, there is a clear mismatch between the battery supporting ability and the diverse power outage events.

e-Energy '17, Shatin, Hong Kong

DOI: http://dx.doi.org/10.1145/3077839.3077863

Table 1: Statistics on number of battery groups for more than 4200 base stations.

number	1	2	3	4	> 4
percent	65.9%	27.5%	4.7%	1.6%	0.1%

Based on the logs of batteries, we further identify the impact of power outages on the conditions of the battery groups, and estimate the battery lifetime and reserve time (indicating the duration a battery group can support). Most researches [7, 10, 11, 16, 17, 19] have been proposed to estimate the battery state and lifetime based on the electrochemical theory of batteries, while some recent works [12, 13] are based on large scale data analysis. Different from those works, we build up a deep learning based model considering the impacts among multiple battery groups. We accordingly develop *BatAlloc*, a battery allocation framework that allocates proper numbers of battery groups to each base stations to address the mismatch between the battery supporting ability and the diverse power outage incidents. We present an effective solution that minimizes both the service interruption time and the overall cost. Our trace-driven experiments show that BatAlloc reduces the average service interruption time from 5 hours to almost zero (i.e., nearly full service availability) with only 88% of the overall cost, as compared to the current real deployment.

2 BACKGROUND AND DATA ANALYSIS

In this section, we discuss the collected dataset from a branch of China Mobile and the related observation on the base stations and backup battery groups. The dataset consists of more than 1.5 million records on battery activities, including such information as the base station locations, battery voltages and event records (e.g., power outage, low voltage alert, high voltage alert, etc.), which are used to analyze the current situation of base stations.

2.1 Base Station Power Supply

In practice, base stations use 48v (24 2v cells installed in series) lead acid battery groups as backup power. The rated capacity of a battery group is usually 500AH and it can support about 10-12 hours (i.e., the battery reserve time is 10-12 hours). We observe the number of battery groups from more than 4200 base stations and show it in Tab. 1. We find that about 93.4% of base stations are equipped with one or two battery groups while only very few base stations have more.

The electrical grid in many regions is not always reliable and power outages occur sometimes. For instance, grid transmission lines can be cut off in case of extreme weather (e.g., storm, hurricane and heavy snow). Then the monitoring system in base stations will report the power outage to the maintenance center and an emergent maintenance should be scheduled according to priorities of different base stations. It usually takes a long time for maintenance engineers to take diesel generators as well as other necessary devices to the corresponding base stations in rural places or very remote regions, especially during a severe weather. So the power recover time is quite uncertain and can not be guaranteed.

Fig. 1 extracts power outage situations of all the base stations and shows the statistics of power outages, from which we can find

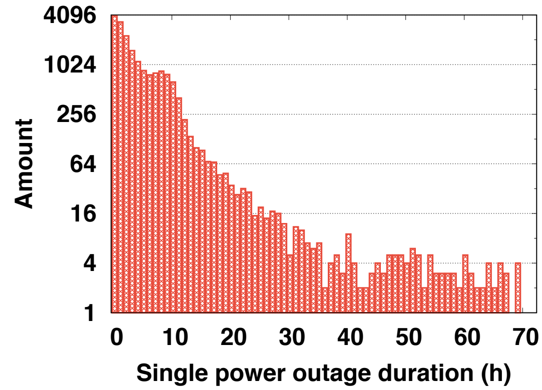

Figure 1: Statistics of power outage duration each time for all base stations.

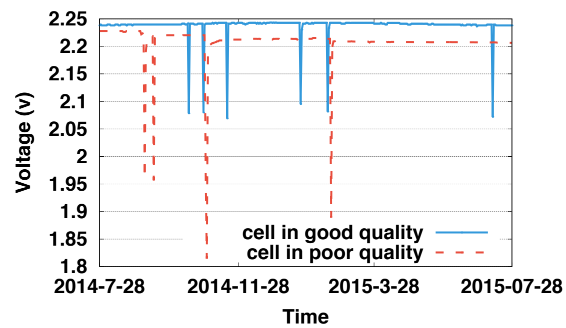

Figure 2: The comparison of two battery cells under different discharge situations.

that quite a few power outages last very long time. However, according to the current battery allocation in Tab. 1, base stations with inefficient backup batteries are not able to sustain the long time power outage without timely emergent maintenance, which can lead to service interruptions and cause serious consequences.

2.2 Backup Battery Activities

Batteries are connected to the electrical grid and kept in float charging state to compensate the capacity loss due to the slow self-discharging process. When there is a power outage, the backup batteries begin to discharge to support base station services. The battery discharging process can be divided into three regions: the coup-de-fouet region [18], the linear region [19] and the hyperbolic region [15] (see details in Appendix A.1).

During a long power outage, the backup batteries may need to discharge to a deep level , which furthers exert an impact on the battery conditions. Fig. 2 presents a comparison of the voltage change between two battery cells that one was in good condition and the other suffered from several deep discharges. We can see that the cell suffered from deep discharges degrades quickly with the float voltage showing a clear decreasing trend where the fast battery degradation contributes to a high battery replacement cost.

In base station power management, a low voltage disconnect (LVD) strategy is applied for battery protection. Base stations have a low LVD settings to prolong the backup power supply, yet actually the deep discharge before LVD has already exerted an impact on battery degradation process (see details in Appendix A.2).

3 BATALLOC FRAMEWORK

Our real trace-driven data analysis clearly reveals that in the battery allocation strategy currently used in practice, there exists a mismatch between the supporting ability of backup batteries and the power outage situations in each base station. The mismatch can lead to serious problems in base stations, such as poor service guarantee, fast battery degradation and high maintenance cost.

One solution is to allocate as many battery groups as possible for every base station, yet such an overprovision will cause a large waste of resources and dramatically increase the overall cost. To this end, we propose BatAlloc, a battery allocation framework to carefully address this mismatch by allocating an appropriate amount of backup battery groups for each base station.

Our BatAlloc framework consists of three major stages. First, we extract the features of base stations from massive data, including the practical distribution of base stations, numbers of battery groups equipped in base stations, power outage situations, etc. Second, we conducts a solid analysis on the battery features, so that battery capacity, battery lifetime and battery degradation under different levels of discharges can be accurately estimated. We develop a deep learning based approach to well model the complicated relationships between different real world events and various battery conditions, which will serve as a key component for the battery allocation optimization. At last, based on the feature profiling results of previous two stages, the battery allocation can then be formulated as an optimization problem. This problem involves multiple optimization goals, e.g., to minimize the service interruptions and minimize the overall cost.

4 BATTERY ALLOCATION SOLUTIONS

In this section, we present how we formulate this battery allocation problem and solve it effectively. A list of notations can be found in Appendix A.3.

4.1 Problem Formulation

Current base stations are mostly equipped with one or two battery groups, which are often insufficient to provide uninterrupted backup power during a long power outage. Assume that we assign n_s battery groups for a particular base station $s \in \mathcal{N}$, where $\mathcal{N}$ is the set of all base stations. We then need to calculate how long the n_s battery groups can support this base station during a power outage. To protect the battery, we disconnect it from the workload when the battery discharges to the end of linear region. To this end, we denote r^t_{s,n_s} as the total reserve time for station s with n_s battery groups at time t.

We denote the time duration from the beginning of power outage to electrical grid recovery or diesel generator launch in station s as $\mathbf{o}_s = \{o_s^{t_1}, o_s^{t_2} \cdots o_s^{t_i}\}$, where t_i is a time-based index. Once the duration exceeds the battery reserve time, there will be a service interruption. We assign importance factor ω_s to represent the service interruption severity (e.g., the service interruptions in core station have more serious consequences). Thus, we have our first optimization objective, which minimizes the total service interruption time:

$$Min: \mathcal{I} = \sum_{s \in \mathcal{N}} \mathcal{I}_s = \sum_{s \in \mathcal{N}} \frac{\omega_s \sum_{t \in \mathcal{T}} [max(0, o_s^t - r^t_{s,n_s})]}{T_{s,n_s}} \quad (1)$$

where $\mathcal{T}$ is the time-based index range of the considered period. We use T_{s,n_s} as denominator for normalization(e.g., representing annual service interruption time).

Besides achieving as short service interruption time as possible, telecom operators may also want to reduce the overall cost, which includes the battery replacement cost and emergent maintenance cost. Battery degradation contributes to battery replacement cost. Besides, when there is a long power outage that the battery capacity is not sufficient enough, engineers may be scheduled an emergent maintenance to the corresponding base station for power generation. We denote c_b as the replacement cost of a single battery group of base stations and $c_{m,s}$ as the emergent maintenance cost of station s. Then we have our second optimization objective, i.e., minimizing the overall cost C_{all} for telecom operators:

$$\begin{aligned} Min: C_{all} &= C_b + C_m \\ &= \sum_{s \in \mathcal{N}} \frac{n_s c_b + \sum_{t \in \mathcal{T}} (x_s^t c_{m,s})}{T_{s,n_s}} \end{aligned} \quad (2)$$

where x_s^t is a binary variable indicating whether there is an emergent maintenance for a power outage.

In practice, there may be other requirements that limit the number of battery groups being installed at a base station. Then we have the following constraint:

$$\forall s, n_L \leq n_s \leq n_U, n_s \in \mathbb{N}^+ \quad (3)$$

Telecom operators usually want to control the overall cost within a give upper budget limit $\mathcal{B}$. So we also have:

$$C_{all} \leq \mathcal{B} \quad (4)$$

4.2 Deep Learning Based Battery Profiling

In order to solve the optimization problem on battery allocation, we first need to model the lifetime and reserve time of the batteries in a base station. We develop a deep learning based approach that utilizes the deep neural network (DNN) to accurately model the voltage trend based on historical events and voltages with the consideration of multiple battery groups.

We first filter out the noise voltage data generated during battery activities (e.g., charging and discharging) and only extract the effective float voltage data. Given a time series of float voltages for battery i, we divide them into a number of time segments where the length of each segment is l. For each segment k, we fit the voltage decreasing trend by linear regression and obtain the voltage change slope s_i^k as well as the initial voltage value v_{if}^k. Then each time segment can be represented as $\{(v_{if}^1, s_i^1), (v_{if}^2, s_i^2) \cdots (v_{if}^k, s_i^k)\}$. We define voltage degradation term as the rate of change on voltage slope for a battery. Then we have battery *degradation* d_i^k as following:

$$d_i^k = s_i^k - s_i^{k-1} \quad (5)$$

For each segment, the battery voltage degradation is ascribed to the battery activities, which are directly reflected by the event logs. We define $\mathbf{e_i}^k = \{e^k_{i,1}, e^k_{i,2}, \cdots e^k_{i,m}\}$ as the input events for battery i in time segment k, where m is the number of event categories. When a base station is equipped with multiple battery groups, the impact of activities is actually shared by all these batteries. Then the impact on each single battery should be proportionally reduced. Thus, we can build up a learning model from events $\frac{\mathbf{e_i}^k}{n_s}$ to the battery degradation d^k_i in segment k, where n_s is the number of battery groups in base station s.

Formally, the inputs are the event sets associated with related segments. Let $\mathcal{E}$ denote the input space of the historical events and we have $\mathcal{E} = \{\frac{\mathbf{e_1}}{n_1}, \frac{\mathbf{e_2}}{n_2}, \cdots \frac{\mathbf{e_N}}{n_N}\}$ with N examples. The outputs are voltage degradations for each segment. Let $\mathcal{D}$ denote the output space of voltage degradation and we have $\mathcal{D} = \{d_1, d_2, \cdots d_N\}$. The modeling process is actually a mapping from $\mathcal{E}$ to $\mathcal{D}$.

We build up a DNN to model the battery degradation process. Each node in the input layer is associated with one kind of events and output layer has one node for degradation estimation. Assuming the base station's situations keep statistically consistent every year, we can then obtain the voltage degradation utilizing our deep learning model. Given that the target time t falls in segment $k+1$ and $v^{t_1}_{if}$ is the initial voltage value, the float voltage can be calculated as follows:

$$v^t_{if} = v^{t_1}_{if} + \sum_{j=1}^{k}(d^j_i + s^{j-1}_i)l + (d^{k+1}_i + s^k_i)(t - kl) \quad (6)$$

With the domain knowledge, a battery is judged in poor quality when its float voltage is below a pre-defined threshold θ. Then we can obtain the lifetime of battery i in station s if the float voltage falls below θ at segment $k+1$:

$$T_i = \frac{\theta - v^{t_1}_{if} - \sum_{j=1}^{k}(d^j_i + s^{j-1}_i)l}{d^{k+1}_i + s^k_i} + kl \quad (7)$$

When there is a power outage, the batteries begin to discharge to provide backup power. According to the electrochemistry knowledge of base station battery features [20], there is a voltage drop from the float charge state to the plateau discharging state mostly due to the cell internal resistance and polarization. We denote the voltage drop as ϵ and we can calculate the plateau discharging voltage as $v^t_i = v^t_{if} - \epsilon$.

Let τ denote the rated battery reserve time of a new battery before the end voltage. The reserve time r^t_{s,n_s} can be calculated as follows (where the estimation of Φ^t_s is described in Appendix A.4):

$$r^t_{s,n_s} = \tau n_s \Phi^t_s \quad (8)$$

4.3 Battery Allocation Algorithm

With the profiling results of base station features and battery features, we now analyze the characteristics of this optimization problem. Intuitively, given the same external incidents happening to a base station, the base station can sustain longer power outages when equipped with more battery groups. The total service interruption time is thus reduced. Meanwhile, since the emergent maintenance is accompanied with service interruptions, fewer service interruptions also cut down the cost of emergent maintenance. Thus in our allocation model when the battery group number keeps increasing, both the service interruption time and the emergent maintenance cost will monotonously decrease until no service interruption occurs.

However, the battery replacement cost is different, where the process can be divided into two stages: In the first stage, when the battery group number of a base station increases, the additional backup power helps the base station sustain long power outages and reduce deep discharging of batteries. In the second stage, if we continue to increase the battery group number, the extra backup power becomes redundant due to enough power supply. Then the service interruption time remains unchanged or decreases very little, while the battery replacement cost increases due to the unavoidable battery degradation process.

Therefore, there are three situations of the battery replacement cost depending on the different conditions of the corresponding base station, i.e., the cost first drops and then rises (both in stage 1 and stage 2), the cost keeps decreasing (only in stage 1), and the cost keeps increasing (only in stage 2). As the sum of battery replacement cost and emergent maintenance cost, the overall cost can also have these characteristics when the battery replacement cost dominates, which is further verified by our real data-driven experiments in §5.

Based on the above analysis, it is easy to see that the two objectives in our model are conflicting and multiple Pareto optimal solutions may exist. Considering the practical situation of telecommunication industry, the most important objective for telecom operators is to provide more reliable cellular communication services. So we utilize a lexicographic method [14] to rank the multiple objectives in the order of importance. We first consider minimizing the service interruption time when the overall cost has an upper limit $\mathcal{B}$. Then we strive to minimize the overall cost without increasing the service interruption time.

We design out heuristic algorithm to divide the solving process into two stages. In the first stage, for each base station we keep increasing the battery group number until the overall cost begins to rise. We thus stop and record the battery allocation results in the first stage.

In the second stage, the two objectives are conflicting because the battery replacement cost begins to rise. As aforementioned, we consider reducing the service interruption time when the overall cost does not exceed the budget limit. To better balance the tradeoff between them, we define *Gain* as the ratio of the weighted service interruption decrease and the overall cost increase:

$$Gain = \frac{\mathcal{I}_s(n_s) - \mathcal{I}_s(n_s+1)}{C_{s,all}(n_s+1) - C_{s,all}(n_s)} \quad (9)$$

We each time select the base station with the maximum *Gain* and add one battery group to it until we reach the budget. By utilizing such a greedy approach we guarantee to reduce the most service interruption time with the least cost increase for each step.

5 EVALUATION

In this section, we present the evaluation of our BatAlloc framework based on real trace-driven experiments. We first evaluate

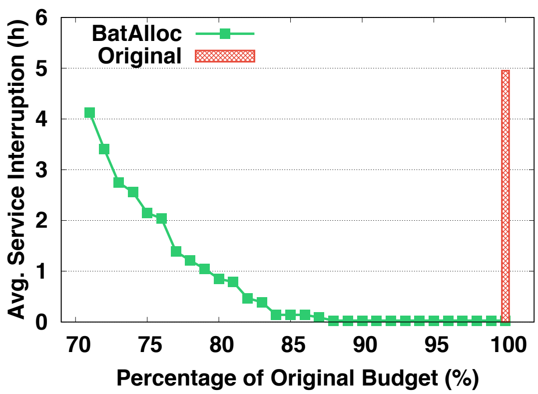

Figure 3: The comparison of overall service interruption.

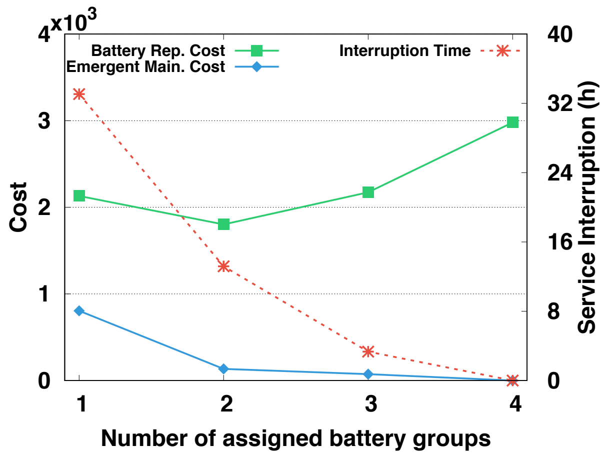

Figure 4: Various metrics for a typical base station.

our battery feature profiling process and compare the predicted results with ground truth. Based on the base station and battery profiling results, we present the performance evaluation on the overall BatAlloc framework.

5.1 Experiment Setup

We conduct data processing on our dataset (about 320GB) from China Mobile and extract useful features on base stations and backup batteries. We construct the deep learning based model using Keras [2], which is a neural network library on top of TensorFlow [5] and Theano [6].

The parameter settings of our experiments are extracted from our dataset as well as adopted from the typical settings based on the domain knowledge. The normal float voltage is 2.25v and the plateau discharging voltage v_P is set as 2.08v for a new battery cell. In our experiments, we set the importance factor ω_s based on the population that a base station covers, which is normalized to a value between 0 and 1. The rest of parameters are set according to the real world market [3, 4].

5.2 Experiment Results

We first evaluate the performance of our deep learning based battery profiling model. We use the data of the first 240 days as the training set and the data of the next 120 days as the testing set. The results shows the RMS error of our deep learning based predicting model is less than 0.01v compared with ground truth. With the domain knowledge of battery features, we can obtain the battery lifetime and reserve time used for the battery allocation optimization in the BatAlloc framework, which will be evaluated next.

For comparison, we extract the current battery deployment as a baseline from the real world dataset and use the Original allocation to represent it. Fig. 3 plots the annual average service interruption time with different budget limit $\mathcal{B}$, where, for ease of comparison, the budget limit is normalized by the baseline budget (i.e., 100% means the budget limit is equal to 100% of the original baseline budget). The service interruption time drops observably as the budget limit increases and we can achieve nearly full service guarantee with only 88% of the baseline budget. These results demonstrate that our BatAlloc framework is capable of providing much more reliable service with a remarkably reduced cost.

We conduct a case study in Fig. 4 when a station is equipped with different numbers of battery groups. As the number of battery groups increases from 1 to 4, the emergent maintenance cost and the service interruption time decrease monotonously due to more sufficient backup power. On the other hand, the battery replacement cost achieves minimum when the number of battery groups is 2 since the additional battery group can drastically reduce the impact of overdischarging and prolong the battery lifetime. If we keep increasing battery groups, the extra battery power continues to reduce the service interruption time, while the battery replacement cost rises largely mostly due to the unavoidable battery degradation.

We put some more intermediate results of experiments in Appendix A.5.

6 CONCLUSION

Current cellular communication base stations are facing serious problems due to the mismatch between the power outage situations and the backup battery supporting abilities. In this paper, we proposed BatAlloc, a battery allocation framework to address this issue. We first conducted a systematical analysis of a massive dataset of base stations and batteries. Then we built up a deep learning based model to precisely capture the battery conditions and further profile the battery features. With the profiling results, we formulated this battery allocation issue as a multi-objective optimization problem and designed an efficient algorithm to solve it. Our real trace-driven experiments showed that compared to the current practical deployment, our framework can remarkably reduce the service interruptions as well as the overall cost.

ACKNOWLEDGMENTS

This work is supported in part by a Canada NSERC Discovery Grant, and an NSERC E.W.R. Steacie Memorial Fellowship. Feng's work is partly supported by the Start-up Grant from the University of Mississippi and NSF I/UCRC 1539990.

REFERENCES

[1] 2010. Power outage in southwestern Connecticut. http://www.nj.com/news/index.ssf/2010/03/storm_brings_power_outages_roa.html. (2010). Accessed: 2017-01-20.

[2] 2015. Keras on github. https://github.com/fchollet/keras. (2015). Accessed: 2017-01-20.

[3] 2016/1/19. Powering Your Cell Towers. http://www.westell.com/assets/MKTG-rA-Powering-Your-Cell-Tower_White-Paper.pdf. (2016/1/19). Accessed: 2017-01-20.

[4] 2017. Labor cost on stationary engineers. http://www.payscale.com/research/CA/Job=Stationary_Engineer/Hourly_Rate. (2017). Accessed: 2017-01-20.

[5] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, and et al. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. (2015). http://tensorflow.org/ Software available from tensorflow.org.

[6] Rami Al-Rfou, Guillaume Alain, Amjad Almahairi, Christof Angermueller, and et al. 2016. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints* abs/1605.02688 (May 2016). http://arxiv.org/abs/1605.02688

[7] Adnan H Anbuky and Phillip E Pascoe. 2000. VRLA battery state-of-charge estimation in telecommunication power systems. *IEEE Transactions on Industrial Electronics* 47, 3 (2000), 565–573.

[8] Adnan H Anbuky, Phillip E Pascoe, and Richard G Lane. 1998. VRLA battery capacity measurement and discharge reserve time prediction. In *Telecommunications Energy Conference, 1998. INTELEC. Twentieth International*. IEEE, 302–310.

[9] Wissam Balshe. 2011. Power system considerations for cell tower applications. *Cummins Power Generation* (2011).

[10] Bikramjit S Bhangu, Paul Bentley, David A Stone, and Christopher M Bingham. 2005. Nonlinear observers for predicting state-of-charge and state-of-health of lead-acid batteries for hybrid-electric vehicles. *IEEE Transactions on Vehicular Technology* 54, 3 (2005), 783–794.

[11] Martin Coleman, Chi Kwan Lee, Chunbo Zhu, and William Gerard Hurley. 2007. State-of-charge determination from EMF voltage estimation: Using impedance, terminal voltage, and current for lead-acid and lithium-ion batteries. *IEEE Transactions on industrial electronics* 54, 5 (2007), 2550–2557.

[12] Xiaoyi Fan, Feng Wang, and Jiangchuan Liu. 2016. Boosting Service Availability for Base Stations of Cellular Networks by Event-driven Battery Profiling. *ACM SIGMETRICS Performance Evaluation Review* 44, 2 (2016), 88–93.

[13] Xiaoyi Fan, Feng Wang, and Jiangchuan Liu. 2016. On Backup Battery Data in Base Stations of Mobile Networks: Measurement, Analysis, and Optimization. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*. ACM, 1513–1522.

[14] C-L Hwang and Abu Syed Md Masud. 2012. *Multiple objective decision making—methods and applications: a state-of-the-art survey*. Vol. 164. Springer Science & Business Media.

[15] I Kurisawa and M Iwate. 1997. Capacity estimating method of lead-acid battery by short-time discharge-introduction of approximate expression of discharge curve obtained by compound expression of straight line and hyperbola. In *Telecommunications Energy Conference, 1997. INTELEC 97., 19th International*. IEEE, 483–490.

[16] Koray Kutluay, Yigit Cadirci, Yakup S Ozkazanc, and Isik Cadirci. 2005. A new online state-of-charge estimation and monitoring system for sealed lead-acid batteries in telecommunication power supplies. *IEEE Transactions on Industrial Electronics* 52, 5 (2005), 1315–1327.

[17] Valentin Muenzel, Julian de Hoog, Marcus Brazil, Arun Vishwanath, and Shivkumar Kalyanaraman. 2015. A Multi-Factor Battery Cycle Life Prediction Methodology for Optimal Battery Management. In *Proceedings of the 2015 ACM Sixth International Conference on Future Energy Systems*. ACM, 57–66.

[18] Phillip E Pascoe and Adnan H Anbuky. 2002. The behaviour of the coup de fouet of valve-regulated lead–acid batteries. *Journal of power sources* 111, 2 (2002), 304–319.

[19] Pillip E Pascoe and Adnan H Anbuky. 2004. VRLA battery discharge reserve time estimation. *IEEE transactions on power electronics* 19, 6 (2004), 1515–1522.

[20] Whitham D Reeve. 2006. *DC power system design for telecommunications*. Vol. 14. John Wiley & Sons.

A APPENDIX

A.1 Battery Discharge Process

Fig. 5 illustrates a typical discharging curve for a lead acid cell. The coup-de-fouet region appears at the start of battery discharging, where the battery voltage first falls quickly below its open circuit voltage and then rises to a higher plateau voltage in a short time. This kind of voltage change is a special characteristic usually observed from lead acid batteries. Then the discharging process goes into a long linear region, where the voltage drop has an approximately linear relationship with the discharging time. The discharging characteristic is robust to variations in operating conditions as well as battery conditions, such as the discharging mode (constant current or constant power), ambient temperature, battery degradation condition [19], etc. A battery will release most of its electrical energy during the linear region. In the last hyperbolic region, the voltage falls very fast while it can only release a very small fraction of power.

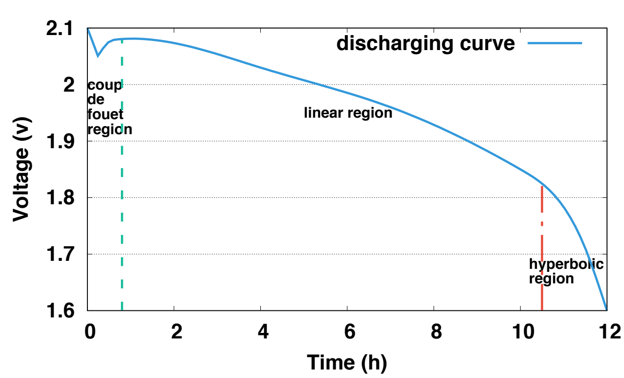

Figure 5: Typical discharge voltage versus time characteristics.

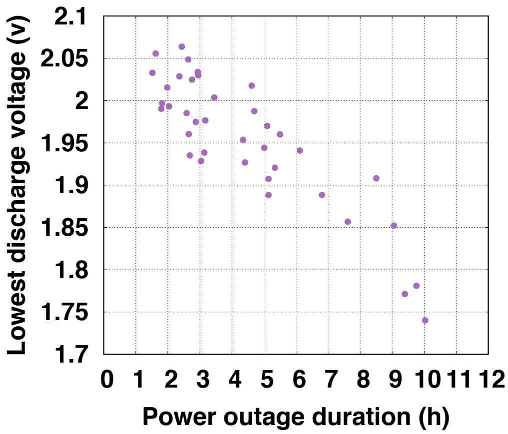

Figure 6: The relationship between power outage duration and voltage drop.

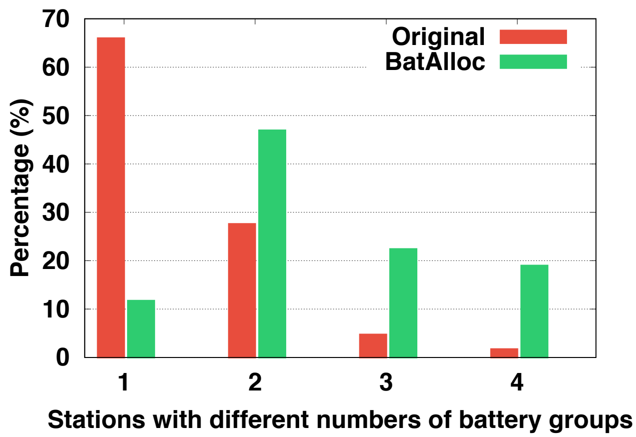

Figure 7: The battery allocation result by BatAlloc.

The conditions of lead acid batteries are largely dependent on the depth of discharge (DoD). If a lead acid battery frequently discharges to high DoDs, the lead in the negative plate will form large lead sulfate crystals adhered to the negative plate and further accelerate the battery sulfation. This degradation process is accumulative, which as a result greatly reduces the capacity and lifetime of lead acid batteries. Therefore, it is not desirable to allow a battery group to discharge completely, because the battery group will be permanently damaged and become incapable of being fully recharged to its rated capacity again. According to the industry standard, a battery should be replaced once its capacity falls below the 80% of the rated capacity.

A.2 Low Voltage Disconnect Strategy

When the battery voltage falls below a first pre-defined threshold, the lead acid battery groups will be disconnected from the secondary devices and only provide backup power to primary communication devices. When the voltage continues to drop below a second predefined threshold, power system cuts off all the loads to avoid the battery groups from being drained. Fig. 6 plots the relationship between the power outage duration and the voltage drop (to eliminate the impact of battery group numbers, we only choose those base stations with one battery group). We observe that the discharge voltage could fall below 1.75v during a long power outage, which in fact will seriously damage the battery condition.

A.3 Notation of BatAlloc Model

Table 2: Notations

n_s	number of battery groups at base station s
r^t_{s,n_s}	the reserve time for station s at t with n_s battery groups
o^t_s	duration from power outage to grid recovery or generator launch for stations s at t
ω_s	the importance factor of station s on service interruption severity
T_{s,n_s}	the expected lifetime of each battery group when stations s is equipped with n_s battery groups
$\mathcal{I}$	the normalized total service interruption time
$\mathcal{T}$	the time-based index range
$\mathcal{N}$	the set of all the base stations
c_b	the replacement cost of a battery group
C_b	normalized total cost on battery group replacement
x^t_s	a variable indicating whether station s needs an emergent maintenance at t
$c_{m,s}$	emergent maintenance cost for station s
C_m	the normalized total emergent maintenance cost
C_{all}	the normalized overall cost
n_L	lower limit of battery group number in a station
n_U	upper limit of battery group number in a station
$\mathcal{B}$	the budget limit
v^k_{if}	float voltage of battery i in k-th segment
s^k_i	voltage slope of battery i in k-th segment
d^k_i	degradation of battery i in k-th segment
$\mathbf{e}^k_i$	the event set for battery i in k-th segment
$\mathcal{E}$	the event set
$\mathcal{D}$	the degradation set
ϵ	the voltage drop at start of discharging
Φ^t_s	percentage of remaining capacity of a battery group
v^t_i	the plateau discharging voltage of battery i at t
v_E	the end discharging voltage in linear region
v_P	plateau voltage at the beginning of discharging
τ	the rated reserve time before end voltage

A.4 State of Charge Estimation

The battery discharge characteristics can be utilized to estimate the battery state of charge (SOC) and battery reserve time [8, 19] in the linear region. The scaled discharge curves of batteries with different degradation keep highly consistent, and the plateau discharge voltage drops with the degradation level. Thus we can build the mapping from the plateau discharge voltage to the corresponding capacity in the linear region. Let v_E denote the end voltage and v_P is the plateau voltage of discharging phase for a new battery cell. We use Φ^t_s to represent the percentage of remaining capacity of a

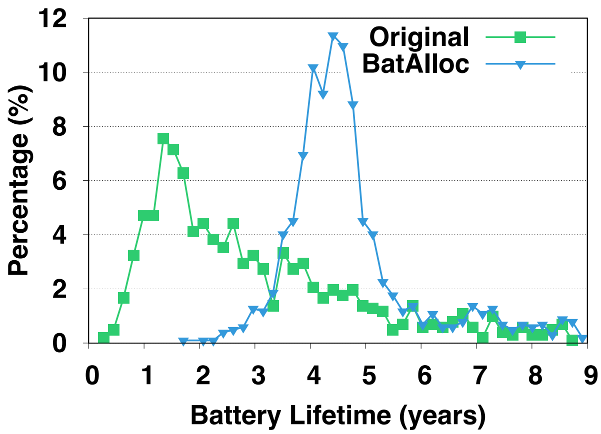

Figure 8: The comparison of battery lifetime.

battery group in the linear region at t. Then we can calculate Φ_s^t based on the discharging voltage v_i^t:

$$\Phi_s^t = \frac{v^t - v_E}{v_P - v_E} \tag{10}$$

A.5 Additional Experiment Results

Fig. 7 compares the different allocation results on battery group number between the original allocation scheme and our BatAlloc framework. The original battery allocation result is largely skewed that over 65% base stations are equipped with only one battery group. Our framework considers both the base station situations and battery features, allocating 2 battery groups to most base stations and 3 or 4 battery groups to those with prolonged power outages.

We also investigate the impact of different battery allocation strategies on battery lifetime. As shown in Fig. 8, in the original allocation the average battery lifetime is only around 1.5 years and far less than expected. After using the BatAlloc to allocate suitable numbers of battery groups for base stations, the average battery lifetime has achieved to 4.5 years, roughly 2 times longer than that of the original allocation. The results indicate that our framework can also better protect base station batteries and significantly prolong their average lifetimes.